REMARKS

Claims 1-10 are pending in the current application. In an office action dated November 27, 2007 ("Office Action"), the Examiner rejected claim 1 under 35 U.S.C. §103(a) as being unpatentable over Yamada et al., U.S. Patent No. 5,455,820 ("Yamada") in view of Homberg et al., U.S. Patent No. 6,661,802 B1 ("Homberg"), rejected claims 2-5 under 35 U.S.C. §103(a) as being unpatentable over Yamada in view of Homberg and further in view of Chiussi et al., U.S. Patent No. 5,701,292 ("Chiussi"), rejected claims 1, 6-9, and 10 under 35 U.S.C. §103(a) as being unpatentable over Morgenstern et al., U.S. Patent No. 6,614,756 B1 ("Morgenstern") in view of Robles et al., U.S. Patent No. 6,282,172 B1 ("Robles"), and rejected claims 2-3 and 4-5 under 35 U.S.C. §103(a) as being unpatentable over Morgenstern in view of Robles and further in view of Chiussi. Applicants' representative respectfully traverses the 35 U.S.C. §103(a) rejections.

The current invention, as claimed in claims 1-10, is not particularly difficult to understand. It is therefore surprising that, again, the Examiner appears to have failed to appreciate various aspects of the currently claimed invention. One embodiment of the current invention is thoroughly and clearly illustrated in Figures 13A-J, and discussed in the current application beginning on line 3 of page 15. Figure 13 shows three receive queues, or input queues, 1302-1304, and three transmit queues, or output queues, 1306-1308. Each queue is shown having at least eight queue entries for at least eight different packets, or messages, although, in practical embodiments, input and output queues may have far larger capacities. Each queue is associated with a communications port including a transceiver. As discussed in the current application, Figures 13A-J are abstractions of a network multiplexer, since, in a network multiplexer, as shown in Figure 7 of the current application, each communications port is associated with both a receive queue and a transmit queue. Each output queue 1306-1308 is associated with a high-threshold value (1314 for output queue 1306) and a low-threshold value (1316 for output queue 1306). Each input queue, or receive queue, is associated with a bit array, in the embodiment discussed with reference to Figures 13A-J, such as bit

array 1318 associated with input queue 1302. Each bit in the bit array represents a different one of the three output queues 1306-1308. When a bit in the bit array is set to Boolean value "1," then the receive queue has received a flow-control directive from the corresponding transmit queue. As clearly summarized in the summary-of-the-invention section of the current application, beginning on line 30 of page 2:

> If the number of queued message descriptors in a transmit queue exceeds the high threshold, any port thereafter attempting to queue additional message descriptors to the transmit queue are directed by the port associated with the transmit queue to ***undertake flow control on their associated network media in order to temporarily prevent reception of additional communications packets.*** Once the number of message descriptors queued to the transmit queue falls below the low threshold, all ports to which the port associated with the transmit queue has sent flow control directives are sent release flow control messages so that these ports can discontinue flow control and resume receiving communications packets. A flow controlled port remains flow controlled until all outstanding flow control directives have been removed by subsequent release flow control messages. (emphasis added)

Thus, flow control refers to blocking a communications medium, by a port attached to the communications medium, from sending further packets to the port until the flow-control is released.

In Figure 13A, each of the input queues 1302-1304 is filled with packets, labeled "A," "B," and "C," for input queues 1302, 1303, and 1304, respectively. Additional, received packets are shown, in Figure 13B, by the capital letters "A," "B," and "C," associated with arrows representing input to input ports, such as arrow 1330. At the point in time represented in Figure 13C, transmit queue 1306 contains ten packets, or messages, and thus has reached the high-threshold value (1314 in Figure 13C) of "10." Therefore, the transmit queue indicates to input queue 1302, from which the last message was received by transmit queue 1306, that it does not wish to receive further messages. Accordingly, as shown in Figure 13D, a "1" value is placed in the bit array 1318 associated with input queue 1302 to indicate that the port associated with input queue 1302 should be flow controlled. As discussed in the current application, once the bit has been set, as shown in Figure 13E, the port is allowed to receive two more messages

before, as shown in Figure 13F, input to the port associated with input queue 1302 is halted. As discussed in the current application, this additional number of received messages following a directive to flow control the port is computed so that the medium to which the port is connected has time to discontinue message transmission, without dropping any already sent messages or messages being transmitted at the time that flow-control is exercised. Flow-control directives are also sent to the input queue 1303, as shown in Figure 13F, since both output queues 1306 and 1307 have exceeded the high-threshold value of "10." When the number of messages queued to output queue 1306 falls below the low-threshold value (1316 in Figure 13G), then the input queues 1302 and 1303 receive flow-resumption directives from the port associated with output queue 1306, as indicated by the clearing of the bits corresponding to output queue 1306 in the bit fields, such as bit-array 1318, associated with input queues 1302 and 1303. After a period of time, message reception resumes through the port associated with input queue 1302, as shown in Figure 13I. When the number of messages queued to output queue 1307 falls below the minimum-threshold value (1340 in Figure 13I), then flow-control is released with respect to input queue 1303, as shown in Figure 13I by the clearing of the bit 1337 associated with output queue 1307, and message reception resumes through the port associated with input queue 1303, as shown in Figure 13J.

In attempting to read the current claims onto Yamada, in the first rejection of claim 1, the Examiner has attempted to redefine the phrase "flow control request" to mean "buffer occupancy signal state signal," as stated by the Examiner in the final four lines of page 3 of the Office Action. In Yamada, as clearly stated by Yamada in the cited passage of lines 40-47 of column 3, each output buffer section includes "a buffer occupancy ratio calculator 320 for calculating an occupancy ratio of the cell buffer and outputting a signal representative of the calculated ratio, i.e., a buffer occupancy state signal, to all input buffer sections when the buffer occupancy ratio has exceeded a predetermined threshold." This is not a flow-control request. As further stated by Yamada, beginning on line 53 of column 3: "In each input buffer section 100, the buffer controller 120 includes control means for using, when the buffer occupancy state signal is absent, only one of the cell buffers 140_1-140_3." However, when an output buffer begins

to overflow, packets or messages, referred to as "cells" in ATM networking, are buffered, first in the output buffer, and then in the input buffer, so that no cells are ever discarded, as described in the paragraph beginning on line 50 of column 4:

> The procedure of FIG. 4(*a*) begins with a step S11 in which the buffer controller 120 sees that a certain output buffer section has overflown in response to the associated buffer occupancy state signal line. Then, for the recovery of the output buffer section from the overflow, the buffer controller 120 interrupts the flow of cells into the output buffer section of interest. At the same time, the buffer controller 120 selects a spare cell buffer which is included in the input buffer section to prevent cells from being discarded. In the illustrative embodiment, the cell buffers 140_2 and 140_3 of each input buffer section are assumed to be spare cell buffers; the buffer controller 120 selects one of time (steps S12-S15). If both the cell buffers 140_2 and 140_3 are full (N, step S13), the buffer controller 120 sends a command to the address filter 110 to prevent it from gating cells addressed to the overflown output buffer section (step S17). ***This would cause such cells to be discarded. However, the input buffer section is provided with a number of cell buffers great enough to avoid such an occurrence.*** (emphasis added)

In other words, there are no flow-control requests made by, or executed by, any component described by Yamada. As clearly defined in the current application, and as well known to anyone familiar with networking systems, a flow-control request requests that transmission of messages through a communications medium to a port be discontinued. In Yamada, there is no discontinuing of receptions of cells by the input buffer. Instead, rather than using a single input buffer associated with each input signal line, Yamada's system employs multiple buffers to buffer incoming cells when an output buffer becomes filled or congested. However, a cursory glance at any of Figures 13A-J, Figure 7, or many other figures of the current application, reveals that, in the described embodiment of the current invention, multiple input and output buffers are always used within each input or output queue of the described network multiplexer. Thus, Yamada simply describes a method for resorting to using multiple input and output buffers, rather than single input and output buffers, under high load. This has nothing to do with flow-control requests. Applicants' representative has failed to find the phrase "flow control" in Yamada.

In the rejection of claim 1, the Examiner states that the absence of an occupancy signal when buffer occupancy falls below a predetermined threshold, in Yamada, constitutes a low threshold. As the Examiner surely appreciates, Yamada mentions only a single threshold value - the predetermined threshold value representing an occupancy ratio above which a buffer occupancy state signal is sent. In other words, when the occupancy is below the predetermined threshold value, in Yamada, no signal is sent, and when the occupancy ratio is above the predetermined threshold value, a signal is sent. No twisting of Yamada's clearly stated disclosure can possibly create two different threshold values associated with each output buffer from a single predetermined threshold value. The absence of a signal does not constitute a positively claimed low-threshold value. Claim interpretation would become utterly arbitrary and meaningless were the absence of a signal construed to teach or suggest a positively-claimed low threshold value. By such reasoning, the absence of a signal could teach anything, from a car transmission to a protein-based pharmaceutical.

In the rejections of claim 1-5, in which Yamada is used as the primary reference, the Examiner states: "Yamada et al. discloses all the claimed limitation with exception of being silent with regard to the following features: regarding claim 1, when the transmit queue currently contains a maximum number of message descriptors, discard the message descriptor." Yamada is not at all silent on this matter. In fact, the entire point of Yamada's disclosure is, as stated in the last sentence of Yamada's abstract: "As a result, when any of the output buffer sections overflows, it is recovered immediately without cells being discarded." At the end of the above-quoted passage of Yamada, Yamada states: "This would cause such cells to be discarded. However, the input buffer section is provided with a number of cell buffers great enough to avoid such an occurrence." It is absolutely and unambiguously clear, from Yamada's statements, that Yamada developed an ATM-switch system to avoid discarding messages. Thus, combining Yamada with an unrelated reference, such as Homberg, that does teach dropping of frames or messages, makes absolutely no sense, and would create a system opposite from both Yamada's disclosed system and Yamada's motivations for implementing Yamada's disclosed system.

In short, the rejections based on Yamada are unfounded and the Examiner's statements incorrect. Yamada is directed to a system that avoids the need for flow control, while the currently claimed invention employs flow control. Yamada fails to teach or suggest both a high threshold value and a low threshold value associated with each transmit queue.

The additional rejections based on Morgenstern and Robles fail, for reasons similar to the failure of the rejections based on Yamada. For example, the Examiner attempts to read the phrase "flow control" onto "detecting signaling congestion situation" recited in lines 23-36 of column 4 of Morgenstern. Lines 23-36 of column 4 of Morgenstern read, as follows:

> The present invention is method of detecting a signaling congestion situation in a transmitter within a switch and for handling and recovering from the signaling congestion. The invention also comprises a method for detecting the absence of a signaling congestion situation and the processing thereof. The invention is applicable to ATM switching networks wherein a sliding window technique is used in transmitting signaling or any other type of messages from a source to a destination. The invention, however, is not limited to application only to ATM networks. It is applicable to any type of communications system whereby a siding window technique is used to transmit data from one point to another.

As can be appreciated by anyone with even cursory understanding of technical matters, the cited passage of Morgenstern does not once mention flow control. Detecting congestion means exactly what it states — namely, detecting when there is insufficient resources to buffer incoming messages prior to transmission. As discussed above, flow control refers to halting message transmission through a communications medium to a port attached to the communications medium. Nothing in the above-quoted passage is in any way related to flow control.

The Examiner cites memory 48 in Figure 3 of Morgenstern as teaching a transmit queue associated with a transmitting port. Note that, in Figure 7 of the current application, each port 702-707 in a network multiplexer is associated with both a receive queue and a transmit queue (708 and 710 for port 702). By contrast, in Morgenstern, no port is associated with memory 48. Instead, memory 48 serves as a centralized buffer

pool managed by controller 50, as stated in the cited passage of Morgenstern, as follows: "A portion of the memory 48 is designated for use as a centralized buffer pool for signaling messages (PDUs) and is of size M."

The Examiner next reads the claim phrase "sending a flow control request" onto "declaring a port to be in a congested state," as recited in Morgenstern in a passage beginning on line 60 of column 4 and running to line 9 of column 5 and in a passage on lines 51-58 of column 9. However, neither of these passages in any way suggests any kind of flow control. The first passage simply states that, in a switch or other communications system implemented according to Morgenstern's disclosure, input ports do not route cells toward an output port that is in a congestion state. Similarly, the cited portion of column 9 mentions nothing about flow control. Morgenstern appears to be stating, in these passages, that an input port stops routing cells to a congested output port. In fact, beginning on line 46 of column 5, Morgenstern states that the input port may then attempt to route a newly arrived cell to a different output port that is not in a congestion state. This is not flow control. As discussed above, flow control involves an input port notifying external devices attached to communications medium that they need to stop sending messages to the input port. Flow control is well understood in networking. The cited passages of Morgenstern do not in any way discuss or suggest flow control. The input and output ports are all contained in a single switch of other communications device. The disclosed system reroutes cells within the switch, but does not undertake flow control.

The Examiner attempts to read the claim language "providing each transmitting port in the network multiplexer with a high threshold and a low threshold" onto lines 10-20 and lines 20-34 of column 5 of Morgenstern. However, a careful reading of these passages reveals that the first threshold is indeed associated with each transmit queue, but that the second threshold is associated with a memory buffer pool; the memory 48 component shown in Figure 3 of Morgenstern, that is not associated with any particular output port or output queue. Again, Morgenstern does not teach or disclose that for which the Examiner cites Morgenstern.

Claim 6 recites "a number of ports, each port having a transceiver and a communications controller." As can clearly be seen in Figure 3 of Morgenstern, Morgenstern's device has input ports that are receivers and output ports that are transmitters, "Tx" means, to those even cursorily familiar with technology, "transmitter," and "Rx" means "receiver." Figure 3 of Morgenstern does not show transceivers. Neither does the cited portion of column 7, on lines 41-57. The Examiner cites cols 60-66 on line 4 of page 11 of the Office Action, but there are no such columns in Morgenstern. The Examiner cites controller 50 of Morgenstern apparently as a controller which each port has, but cursory examination of Figure 3 reveals that controller 50 controls all of the ports in Morgenstern's device, and that no port in Morgenstern's has a controller. Thus, Morgenstern fails to teach or suggest "a receive queue and a transmit queue associated with each port." This again points to the Examiner's failure to actually find teachings or suggestions of claim elements. Instead, the Examiner points to passages unrelated to the claim language.

Claim 6 recites "a receive queue and a transmit queue associated with each port." As can clearly be seen in Figure 3 of Morgenstern, Morgenstern's device has a transmit queue associated with each output line and no queues associated with input ports. The Examiner cites lines 41-57 of column 7 as teaching an input port with memory - but the cited passage only teaches that the input ports are coupled to a data source - a signal line, as shown in Figure 3. Furthermore, a memory is not necessarily a queue. Thus, Morgenstern completely fails to teach or suggest "a receive queue and a transmit queue associated with each port." This again points to the Examiner's failure to actually find teachings or suggestions of claim elements. Instead, the Examiner points to passages unrelated to the claim language.

Again, as in the first set of rejections, the Examiner notes that Morgenstern does not disclose dropping of messages when transmit queues are full, and therefore cites Robles for that claim element. However, again, the Examiner has completely failed to appreciate the intent of Morgenstern's disclosure. Apparently, the Examiner is not familiar with ATM networking protocols. Yamada and Morgenstern are both concerned with ATM networking. As anyone even remotely familiar with

networking protocols well understands, ATM networking involves virtual or logical circuits. A logical circuit is established between a transmitting port and a receiving port prior to the transmission and reception of any data packets by the two ports. Bandwidth is guaranteed, in advance, in ATM networks. Therefore, dropping packets or cells, as they are called in ATM networking, is not a desirable method for dealing with overflow conditions. That is why Yamada and Morgenstern do not teach, mention, or suggest dropping cells as a strategy for handling overflow conditions. Robles is directed to the Internet protocol, a store-and-forward packet-based system that does not involve virtual circuits and guaranteed bandwidth. Such protocols, unlike ATM, are designed to use message dropping for handling overflow conditions.

All of the Examiner's rejections rely principally on either Yamada or Morgenstern. Neither reference teaches, mentions, or suggests that for which the Examiner cites the reference. The Examiner has made no attempt to provide rejections based on any of the exemplary rationales, listed in M.P.E.P. § 2143, other than rationale (G), in which an examiner relies on a teaching, suggestion, or motivation that would lead one of ordinary skill to modify a prior art reference or combine prior art references to arrive at a claimed invention. However, the Examiner has failed to find teachings, suggestions, or motivations for many of the claim limitations, as discussed above. As one example, the Examiner has failed to recite any teaching or suggestion of flow control in Yamada or Morgenstern, despite the Examiner's attempts to read the claim phrase "flow control" onto unrelated passages of Yamada and Morgenstern. The Examiner's rejections are filled with misstatements and inaccuracies, only some of which are listed above.

Applicants' representative would be quite happy to discuss the current application with the Examiner, by telephone, should the Examiner wish to try to better understand the subject matter of the current claims. Applicants' representative would request, however, the Examiner more carefully review references before citing unrelated passages as teaching claim elements which they fail to mention or suggest in any way.

In Applicants' representative's opinion, all of the claims remaining in the current application are clearly allowable. Favorable consideration and a Notice of Allowance are earnestly solicited.

Respectfully submitted,
Bruce W. Melvin et al.
Olympic Patent Works PLLC

Robert W. Bergstrom
Registration No. 39,906

Enclosures:
Postcards (2)
Transmittal in duplicate

Olympic Patent Works PLLC
P.O. Box 4277
Seattle, WA 98194-0277
206.621.1933 telephone
206.621.5302 fax